

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 6

REMARKS

Claim objections

Claim 7 has been objected to because the phrase "component each comprise" should be "component each comprises". Applicant has amended claim 7 in this respect.

Claim rejections under 35 USC 102

Claims 1-4 and 7-20 have been rejected under 35 USC 102(e) as being anticipated by Albert (6,704,278). Applicant respectfully traverses this rejection.

Claims 1-4 and 7-10

Claim 1 is an independent claim, from which claims 2-4 and 7-10 ultimately depend. Applicant submits that claim 1 is not anticipated by Albert, such that claims 2-4 and 7-10 are patentable for at least the same reasons.

Applicant notes first that claim 1 covers two different types of node failover. In one type, the manager component "selects one of the alternate routes to route the destination address to a second node" upon failure of a first node. In another type, switch port remapping is employed. Two different kinds of switch port remapping are covered by claim 1. A first switch "remaps a destination address initially mapped to the port for the third node to the port for the fourth node" upon failure of the third node. A second switch "uses the expanded port range to remap a destination address initially mapped to the input port for the fifth node to the input port for the sixth node" upon failure of the fifth node. Applicant notes that such limitations in claim 1 are consistent with the title of the patent application, "network node failover using path rerouting by manager component or switch port remapping."

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 7

Albert thus has to disclose both of these different types of node failover – path rerouting and switch port remapping – in order to anticipate the claimed invention. Furthermore, Albert has to disclose both kinds of switch port remapping, “remapping a destination address initially mapped to the port for the third node to the port for the fourth node,” as is accomplished by the first switch in claim 1, and “us[ing] the expanded port range to remap a destination address initially mapped to the input port for the fifth node to the input port for the sixth node,” as is accomplished by the second switch in claim 1, in order for Albert to anticipate claim 1.

As a first matter, Applicant strongly contends that Albert does not disclose network node failover – i.e., the failure of a node and the remapping or rerouting of that node to another node – at all, in contradistinction to the basic subject matter of claim 1. As is evidenced by the title of Albert, “stateful failover of *service managers*,” Albert pertains to the failure of a primary service manager and the alternate usage of a backup service manager. But claim 1 does not deal with the failover of the manager component, but rather the failure and failover of a node, like a server. The summary of the invention in Albert states that it discloses “[a] system that includes a primary service manager and a backup service manager.” (Col. 3, ll. 50-51) Albert nowhere discloses a primary node and a backup node, and the claimed invention is concerned with the failure and failover of nodes, not the failure and failover of a manager component.

In rejecting claim 1, the Examiner continually relies upon column 11, lines 5-11 of Albert in disclosing the failure of a node. However, column 11, lines 5-11 of Albert read as follows:

As an example, client 304 may wish to establish a TCP connection with a virtual machine having a virtual IP address. It should be noted that other types of connections may also be established. To establish the TCP connection, client 304 sends a SYN packet with a destination address corresponding to the virtual IP address. The SYN packet is received by the forwarding agent 302.

Column 11, lines 5-11 of Albert thus have nothing to do with the failure of a node like a server. Applicant has thoroughly reviewed Albert, and cannot find any disclosure within Albert relating to the failure and failover of a node like a server. Rather, as has been noted, Albert completely pertains to the failure and failure of a service manager. Because claim 1 does not pertain at all to

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 8

the failure and failover of its manager component, but rather to the failure and failover of nodes like servers, Albert in this most basic way cannot anticipate claim 1.

Applicant next goes through the three elements of claim 1 to show how Albert does not disclose the two types of node failure and failover to which claim 1 is limited. The first element of claim 1 is:

a manager component of a network having programmed therein alternate routes for a destination address, *such that upon failure of a first node of the network to which the destination address is initially routed, the manager component selects one of the alternate routes to route the destination address to a second node of the network.*

In relation to the italicized text above, as has been described, Albert already does not anticipate the first element of claim 1, because it does not pertain to or disclose the failure of a first node, like the server 221 of FIG. 2A of Albert. Column 5, lines 5-11 of Albert are relied upon to disclose the failure of such a first node, but as has been described above, these lines of Albert have nothing to do with failure of a node or other server. As has also been described, Albert pertains to the failure of the manager component itself, to which claim 1 is not directed.

Next is the question of whether Albert discloses a manager component “having programmed therein alternate *routes* for a destination address.” The Examiner finds this aspect of claim 1 in the virtual IP address disclosed in column 8, lines 28-29 of Albert. Albert discloses the following in these lines:

The wildcard affinities specify destination IP addresses that correspond to virtual IP addresses of server clusters that are to be load balanced by the service manager.

Note what Albert is disclosing here: the correspondence of real destination IP addresses to virtual IP addresses. Such correspondence means that virtual IP address A may correspond to destination IP addresses B and C, virtual IP address D and E may correspond to destination IP address F, and so on. These are not *routes* for a destination IP address. A route specifies the path to a given node for a given destination address. Applicant refers the Examiner to FIG. 8 of the patent application as filed for an example of two routes. The solid lines 814A, 814B, and

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 9

814C specify one route, through switches 810 and 806, to the first node 802. The dotted lines 816A, 816B, and 816C specify another route, through switches 810 and 808, to the second node 804. A route thus is *more* than just mapping a virtual address to a destination address, as is accomplished in Albert, and specifically describes the path through switches and/or other network components to reach a given node. For instance, in the example of FIG. 8 of the patent application as filed, Albert would only specify the destination address of the first node 802 as corresponding to a given virtual address. Albert does not specify that the route taken to this address is through the switch 810 and then through the switch 806. Albert, for instance, would not care if the route taken is through the switch 810, through the switch 808, and finally through the switch 806 to reach the first node 802. This is because Albert is concerned only with virtual address to destination address correspondence, and not the actual route, as to which the claimed invention is limited.

The Examiner also relies upon column 6, lines 43-60 of Albert to find the manager component having programmed therein alternate routes for a destination address. This part of Albert reads as follows:

A service manager 241 and a second service manager 242 also communicate with the forwarding agents. The service managers provide the decision making capability that is required to provide a network service such as load balancing. The service managers send specific instructions to each of the forwarding agents detailing how certain flows of packets are to be processed. Such packet processing may include simply routing the packet, gathering statistics about the packet, sending the packet to a service manager, sending a notification that the packet has been seen to a service manager, modifying the packet, or using a special method such as tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header. It should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including packet source and destination addresses and port number, and, in some instances, packet data.

Here, too, Albert does not disclose alternate *routes* for a given destination address to reach different nodes, as in the claimed invention. None of the packet processing described in Albert

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 10

pertains to different routes, describing particular paths, for a given destination address. At best, Albert discloses "tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header." However, tunneling and tag switching do not involve specification of a particular route, but rather, again, pertain to simple address to address correspondence – i.e., rather than destination address A actually corresponding to address A, tunneling and tag switching instead can be used to specify that destination address A corresponds to address B or C, and so on. Neither tunneling or tag switching are relevant to specifying the actual path, or route, as in the claimed invention. For these reasons, too, Albert does not anticipate the claimed invention.

With final respect to the first element of claim 1 is the question of whether Albert discloses "the manager component select[ing] one of the alternate routes to route the destination address to a second node of the network" upon failure of the first node to which the destination address is initially routed, to which claim 1 is limited. The Examiner finds this aspect of claim 1 in column 13, lines 2-9 of Albert. This part of Albert reads as follows:

The service manager then directs the forwarding agent to handle the packets in a certain manner by sending fixed affinities to the forwarding agents for each flow and specifying actions to be performed on the packets. In the example shown, *the action involves translating the destination address from the client to a specific host IP address and translating the source IP address in packets from [sic] the host to a virtual IP address*. Other actions may be defined by fixed affinities including translating other IP address . . .

This excerpt of Albert fails to read on the claimed limitation of claim 1 in two ways. First, Albert is not described as performing these actions (i.e., the service manager directing the forwarding agent to handle packets in certain manner) "upon failure of a first node," as to which the claimed invention is limited. That is, claim 1 is limited to the manager component selecting one of the alternate routes to route the destination address to a second node, upon failure of a first node to which the destination address is already routed. Albert, in column 13, lines 2-9, and elsewhere, does not disclose performing *any* kind of action in response to the failure of a node like a server.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 11

Second, whereas the claimed invention is limited to the manager component selecting “one of the alternate routes to route the destination address to a second node,” Albert, in column 13, lines 2-9, does not select an *alternate route*. As has been described, a route is a specific delineation of a path to a node. It is more than a simple address-to-address correspondence. But here, in the italicized portion of column 13, lines 2-9, noted above, Albert simply translates, or corresponds, or substitutes, one address for another. The action “involves translating the destination address from the client to a specific host IP address” and “translating the source IP address in packets [from] the host to a virtual IP address.” Such translation, in other words, means that the destination address is replaced by the specific host IP address, or the source IP address is replaced by a virtual IP address. Translation does not rise to the description of a route – the manner by which a node can be reached through a network. There can be many routes by which the host corresponding to a translated address can be reached within a network. Because the claimed invention is limited to *routes*, and Albert merely discloses address translation, and not the specific routes to reach given hosts, Albert does not anticipate the claimed invention in this respect, too.

In sum, with respect to the first element of claim 1, Albert does not disclose network node failover using path or route rerouting, as to which the first element of claim 1 is limited and directed. Applicant now discusses why Albert does not disclose network node failover in accordance with second element of claim 1. The second element of claim 1 is:

a first switch of the network having a port for each of at least a third and a fourth node of the network, *such that upon failure of the third node*, the first switch remaps a destination address initially mapped to the port for the third node to the port for the fourth node.

The second element of claim 1 is directed to one kind of the second type of node failover to which claim 1 is limited. Specifically, the second element of claim 1 is directed to one kind of switch port remapping, “remapping a destination address initially mapped to the port for the third node to the port for the fourth node.”

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 12

In relation to the italicized text of the second element, as has been described, Albert in the first instance does not anticipate the second element of claim 1, because it does not pertain to or disclosure the failure of a third node, like any of the servers of FIG. 2A of Albert. Column 11, lines 5-11 of Albert are again relied upon to disclose the failure of such a third node, but as has been described above, these lines of Albert have nothing to do with failure of a server or other node. As has also been described, Albert pertains to the failure of the manager component itself, to which claim 1 is not directed.

Next is the question of whether Albert discloses a switch that “remaps a destination address initially mapped to the port for the third node to the port for the fourth node,” upon failure of the third node. The Examiner relies upon column 6, lines 43-60, and column 13, lines 2-9, of Albert in finding this limitation of the claimed invention in Albert. Applicant first notes, however, that these are the exact same excerpts of Albert that were relied upon to read on the first element of claim 1. However, the second element of claim 1 deals with a very different type of node failover – port remapping – than the first element of claim 1 does, which deals with alternate route selection. Therefore, the Examiner cannot have it both ways – either Albert discloses in column 6, lines 43-60, and column 13, lines 2-9 the alternate route selection of the first element of claim 1, or the port remapping of the second element of claim 1. Column 6, lines 43-60, and column 13, lines 2-9 only disclose one type of “thing,” whereas the first and second elements of claim 1 are limited to two very different types of node failover. For this reason alone, Albert cannot anticipate the second element of claim 1. The Examiner cannot use the same description in Albert, which is limited to one type of “thing” to read on two very different types of node failover in the first and second elements of claim 1. To repeat an often-used phrase, one cannot have his cake and eat it, too!

Indeed, Albert, in column 6, lines 43-60, and column 13, lines 2-9, also does not disclose port remapping as is accomplished in the second element of claim 1. First, as has been described above, the actions performed in column 6, lines 43-60, and column 13, lines 2-9, are not

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 13

performed "upon failure of a node," as to which the claimed invention is limited in this second element, too. That is, claim 1 is limited to the first switch remapping a destination address to a different port upon failure of a third node. Albert, in column 6, lines 43-60, and column 13, lines 2-9, as well as elsewhere, does not disclose performing *any* kind of action in response to the failure of a node like a server.

Next, column 6, lines 43-60, of Albert are again repeated to show how port remapping is not accomplished in Albert. These lines of Albert read as follows:

A service manager 241 and a second service manager 242 also communicate with the forwarding agents. The service managers provide the decision making capability that is required to provide a network service such as load balancing. The service managers send specific instructions to each of the forwarding agents detailing how certain flows of packets are to be processed. Such packet processing may include simply routing the packet, gathering statistics about the packet, sending the packet to a service manager, sending a notification that the packet has been seen to a service manager, modifying the packet, or using a special method such as tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header. It should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including packet source and destination addresses and port number, and, in some instances, packet data.

There is only one mention of ports in this entire excerpt of Albert, in the tail end, in which it is parenthetically noted that "it should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including . . . port number." Disclosure that the port number can be modified by a forwarding agent is far cry from, and does not rise to the level of, disclosing "upon failure of a node, remapping a destination address initially mapped to the port for the node to the port for another node," to which the claimed invention is limited. That is, there is nothing in this excerpt of Albert that describes remapping a destination address specifically mapped to one port of one node to another port of another node. The claimed invention is not claiming port number modification by itself, yet this is all that Albert discloses.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 14

Similarly, column 13, lines 2-9 does not disclose port remapping in Albert in the way in which the second element of claim 1 is limited. These lines of Albert are repeated again:

The service manager then directs the forwarding agent to handle the packets in a certain manner by sending fixed affinities to the forwarding agents for each flow and specifying actions to be performed on the packets. In the example shown, the action involves translating the destination address from the client to a specific host IP address and translating the source IP address in packets from [sic] the host to a virtual IP address. Other actions may be defined by fixed affinities including translating other IP address . . .

Ports are not even disclosed at all in this excerpt of Albert. Therefore, it cannot be said that this part of Albert discloses “upon failure of a node, remapping a destination address initially mapped to the port for the node to the port for another node.” It is difficult to imagine how Albert can disclose this aspect of the claimed invention in column 13, lines 2-9, thereof, when the key part of the second element of claim 1 – port remapping – is not even recited. For this reason, too, Albert does not anticipate claim 1.

Next, Applicant discusses the third element of claim 1. The third element of claim 1 is:

a second switch of the network having an input port for each of at least a fifth and a sixth node of the network, and a visible output port and one or more hidden output ports to receive an expanded port range from an assigning manager component, *such that upon failure of the fifth node*, the second switch uses the expanded port range to remap a destination address initially mapped to the input port for the fifth node to the input port for the sixth node.

The third element of claim 1 is directed to another kind of the second type of node failover to which claim 1 is limited. Specifically, the third element of claim 1 is directed to another kind of switch port remapping, “us[ing] the expanded port range [received from an assignment manager component] to remap a destination address initially mapped to the input port for the fifth node to the input port for the sixth node.”

In relation to the italicized text of the third element, as has been described, Albert in the first instance does not anticipate the third element of claim 1, because it does not pertain to or disclose the failure of a fifth node, like any of the servers of FIG. 2A of Albert. The Examiner

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 15

again relies upon column 11, lines 5-11, of Albert to disclose the failure of such a fifth node. However, as has been described above, these lines of Albert have nothing to do with the failure of a server or other node. As has also been described, Albert only pertains to the failure of the manager component itself, to which claim 1 is not directed.

Next, the question is whether Albert discloses a second switch that has “a visible output port and one or more hidden output ports.” The Examiner indicates that the port link from Forwarding Agent 2 to the server 2 in Albert is the visible output port, and that the port link from the Forwarding Agent 2 to the server 2 in Albert is the hidden output port. However, Albert discloses nowhere such port links as being visible or hidden, and in the claimed invention of claim 1, these ports are specifically described and limited to being visible in one instance and hidden in the other instance. Therefore, Albert cannot be said to disclose the claimed invention in this respect, since the generic disclosure of “ports” does not rise to the level of specifically disclosing “visible” and “hidden” ports.

Indeed, the ports being visible and hidden, respectively, are an important part of this aspect of the claimed invention. As described in the patent application as filed,

*The switch may have an expanded port range to allow such remapping *due to it having only one visible port*, such that *one or more other ports are hidden* to a manager component that assigns the switch the expanded port range. *In this way, the second node fails over for the first node.**

(P. 7, para. 36) By not disclosing hidden and visible ports, Albert cannot achieve node failover in the way in which the third element of claim 1 is limited. Albert simply does not disclose hidden as well as visible ports. For these reasons, too, Albert does not anticipate claim 1.

Another question with respect to the third element of claim 1 is whether Albert discloses the second switch “receiv[ing] *an expanded port range* from an assigning manager component.” Here the Examiner relies upon column 6, lines 43-60, and column 13, lines 2-9, of Albert, in finding this limitation of the claimed invention in Albert. Applicant first notes, however, that these are the exact same excerpts of Albert that were relied upon to read on the first element of

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 16

claim 1 and the second element of claim 1. The third element of claim 1 deals with a very different kind of port remapping – the use of an expanded port range received via hidden ports – than the second element of claim 1 (which recites a different kind of port remapping) and the third element of claim 1 (which does not even recite port remapping, but rather path or route rerouting). Again, the Examiner cannot have it both ways – either Albert discloses in column 6, lines 43-60, and column 13, lines 2-9, the alternate route selection of the first element of claim 1, the port remapping of the second element of claim 1, or the different kind of port remapping of the third element of claim 1. Column 6, lines 43-60, and column 13, lines 2-9, *at best or at most* discloses only one type of node failover (and indeed, Applicant contends that these lines of Albert do not disclose any type of node failover whatsoever), whereas the first, second, and third elements of claim 1 are limited to different types and different kinds of node failover. For this reason alone, Albert cannot anticipate the third element of claim 1. The Examiner cannot use the same description in Albert, which is limited to one type of “thing” to read on three very different types and kinds of node failover in the first, second, and third elements of claim 1. To yet again repeat an often-used phrase, one cannot have his cake and eat it, too!

In fact, in column 6, lines 43-60, of Albert, there is no mention of an expanded port range whatsoever. These lines of Albert read as follows:

A service manager 241 and a second service manager 242 also communicate with the forwarding agents. The service managers provide the decision making capability that is required to provide a network service such as load balancing. The service managers send specific instructions to each of the forwarding agents detailing how certain flows of packets are to be processed. Such packet processing may include simply routing the packet, gathering statistics about the packet, sending the packet to a service manager, sending a notification that the packet has been seen to a service manager, modifying the packet, or using a special method such as tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header. It should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including packet source and destination addresses and port number, and, in some instances, packet data.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 17

As has been described, the only disclosure Albert makes as to ports in column 6, lines 43-60 is with respect to "modify[ing] other aspects of packets, including . . . port number." An expanded port range is much more and different than simply modifying port number. You can modify a port number A, for instance, to instead be a port number B. By comparison, an expanded port range presumes that a port number A is within the range of ports B through C. Albert simply discloses the modification of a port number, and not the utilization of an expanded port range. Therefore, column 6, lines 43-60 cannot be said to anticipate the third element of the claimed invention.

Similarly, column 13, lines 2-9 does not disclose expanded port ranges. These lines of Albert are repeated again:

The service manager then directs the forwarding agent to handle the packets in a certain manner by sending fixed affinities to the forwarding agents for each flow and specifying actions to be performed on the packets. In the example shown, the action involves translating the destination address from the client to a specific host IP address and translating the source IP address in packets from [sic] the host to a virtual IP address. Other actions may be defined by fixed affinities including translating other IP address . . .

Ports are not even disclosed at all in this excerpt of Albert. Therefore, it cannot be said that this part of Albert discloses expanded port ranges! It is difficult to imagine how Albert can disclose this aspect of the claimed invention in column 13, lines 2-9, thereof, when a key part of the third element of claim 1 – port ranges – is not even recited. For this reason, too, Albert does not anticipate claim 1.

Finally, Applicant discusses the remaining limitations of the third element of claim 1, in that "the second switch uses the expanded port range to remap a destination address initially mapped to the input port for the fifth node to the input port for the sixth node." Here the Examiner again relies upon column 6, lines 43-60, and column 13, lines 2-9, of Albert, in finding this limitation of the claimed invention in Albert. Applicant first notes, however, that as before these are the exact same excerpts of Albert that were relied upon to read on the first element of claim 1 and the second element of claim 1. The third element of claim 1 deals with a very

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 18

different kind of port remapping – the use of an expanded port range received via hidden ports to achieve port remapping – than the second element of claim 1 (which recites a different kind of port remapping) and the third element of claim 1 (which does not even recite port remapping, but rather path or route rerouting). Again, the Examiner cannot have it both ways – either Albert discloses in column 6, lines 43-60, and column 13, lines 2-9, the alternate route selection of the first element of claim 1, the port remapping of the second element of claim 1, or the different kind of port remapping of the third element of claim 1. Column 6, lines 43-60, and column 13, lines 2-9, *at best* or *at most* discloses only one type of node failover (and indeed, Applicant contends that these lines of Albert do not disclose any type of node failover whatsoever), whereas the first, second, and third elements of claim 1 are limited to different types and different kinds of node failover. For this reason alone, too, Albert cannot anticipate the third element of claim 1. The Examiner cannot use the same description in Albert, which is limited to one type of “thing” to read on three very different types and kinds of node failover in the first, second, and third elements of claim 1.

In the Office Action, the Examiner interprets column 6, lines 43-60, and column 13, lines 2-9, of Albert as indicating that “the Forwarding Agent modifies the initial routing from the server 3 and routes its packet to the server 2 or the sixth node via the port that links the server 2 or the sixth node and its Forwarding Agent 2 in Fig. 2A, col. 6 lines 43-60, col. 13 lines 2-9.” (Page 4 of Office Action) However, even this interpretation does not read on the claimed invention, which is limited to “the second switch us[ing] the expanded port range to remap a destination address.” There is no indication in the Examiner’s summary of column 6, lines 43-60, and column 13, lines 2-9, of Albert that an expanded port range is used at all! Even if the Examiner’s interpretation of these lines of Albert is correct, in other words, it still does not rise to the level of anticipating the claimed invention.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 19

As has been noted, in column 6, lines 43-60, of Albert in particular, there is no mention of an expanded port range whatsoever, and definitely no usage of such an expanded port range to remap a destination address. These lines of Albert read as follows:

A service manager 241 and a second service manager 242 also communicate with the forwarding agents. The service managers provide the decision making capability that is required to provide a network service such as load balancing. The service managers send specific instructions to each of the forwarding agents detailing how certain flows of packets are to be processed. Such packet processing may include simply routing the packet, gathering statistics about the packet, sending the packet to a service manager, sending a notification that the packet has been seen to a service manager, modifying the packet, or using a special method such as tunneling or tag switching to send the packet to a destination other than the destination specified by the destination IP address included in the packet header. It should also be noted that forwarding agents in other embodiments also modify other aspects of packets, including packet source and destination addresses and port number, and, in some instances, packet data.

As has been described, the only disclosure Albert makes as to ports in column 6, lines 43-60 is with respect to "modify[ing] other aspects of packets, including . . . port number." But, how does modifying port number encompass an expanded port range? How does modifying port number equate to remapping a destination address based on an expanded port range? The answer to both of these questions is – modifying port number does not encompass or equate these claimed limitations of the third element of claim 1. Albert simply discloses the modification of a port number, and not the utilization of an expanded port range to remap a destination address. On its face, Albert's parenthetical disclosure of port number modification cannot in any sense be said to rise to the level of remapping a destination address by using an expanded port range. Therefore, column 6, lines 43-60 cannot be said to anticipate the third element of the claimed invention.

Similarly, column 13, lines 2-9 does not disclose utilizing expanded port ranges to remap destination addresses. These lines of Albert are repeated again:

The service manager then directs the forwarding agent to handle the packets in a certain manner by sending fixed affinities to the forwarding agents for each flow and specifying actions to be performed on the packets. In the example shown, the

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 20

action involves translating the destination address from the client to a specific host IP address and translating the source IP address in packets from [sic] the host to a virtual IP address. Other actions may be defined by fixed affinities including translating other IP address . . .

Ports are not even disclosed at all in this excerpt of Albert. Therefore, it cannot be said that this part of Albert discloses usage of expanded port ranges to remap destination addresses. It is difficult to imagine how Albert can disclose this aspect of the claimed invention in column 13, lines 2-9, thereof, when a key part of the third element of claim 1 – port ranges – is not even recited. For this reason, too, Albert does not anticipate claim 1.

Claims 11-15

Claim 11 is an independent claim, from which claims 12-15 depend. Applicant traverses the rejection as to claim 11, such that claims 12-15 are also not anticipated by Albert.

Claim 11 is limited to two different types of node failover. First, one step or act that may be performed for failover is “routing the destination address over an alternate path to the second node selected by a manager component.” Second, another step or act that may be performed for failover is “remapping the destination address from the first port on the switch to a second port on the switch connected to the second node.”

Albert does not disclose either of these two types of node failover. As has been described above in relation to the first element of claim 1, Albert does not disclose rerouting the destination over a different path. As has been described above in relation to the second and third elements of claim 1, Albert does not disclose port remapping. Indeed, as has been described above generally with respect to claim 1, Albert discloses and pertains to only the failover of the manager component itself, not the failover of servers and other types of nodes! For these reasons, Albert does not anticipate claims 11-15.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 21

Claims 16-20

Claim 16 is an independent claim, from which claims 17-20 depend. Applicant traverses the rejection as to claim 16, such that claims 17-20 are also not anticipated by Albert.

Claim 16 is also limited to two different types of node failover. First, the failover action may be "rerouting the destination address to over an alternate path to the failover node from over an original path to the failed node." Second, the failover action may be "remapping the destination address from a first port connected to the failed node to a second port connected to the failover node."

Albert does not disclose either of these two types of node failover. As has been described above in relation to the first element of claim 1, Albert does not disclose rerouting the destination over a different path. As has been described above in relation to the second and third elements of claim 1, Albert does not disclose port remapping. Indeed, as has been described above generally with respect to claim 1, Albert discloses and pertains to only the failover of the manager component itself, not the failover of servers and other types of nodes. For these reasons, Albert does not anticipate claims 16-20.

Claim rejections under 35 USC 103

Claims 5 and 6 have been rejected under 35 USC 103(a) as being unpatentable over Albert in view of Chang (6,724,759). Claims 5 and 6 depend ultimately from claim 1, and therefore are patentable for at least the same reasons that claim 1 is.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 22

Additional pertinent prior art

The Examiner has cited six additional prior art references considered pertinent to Applicant's invention but not specifically relied upon. These references are: Huang (6,308,282); Richter (2003/0097481); Wang (6,757,242); Chen (6,715,098); Blumenau (6,421,711); and, Craft (6,687,758). Applicant has reviewed these references and do not believe that either alone or in combination they render the claimed invention unpatentable. Applicant now discusses each of these references in turn, briefly.

Huang relates to failover a channel. Thus, a primary channel that connects a node to a switch can be failed over to a replacement channel upon the primary channel having a fault. That is, “[f]ailure recovery includes redirecting data transmission of a node from a channel indicating a failure to a stand-by channel.” (Col. 2, ll. 30-32) By comparison, the claimed invention relates to failover of a *node*, not a *channel* connecting to that node, as in Huang. In the claimed invention, a node fails over to another node, by rerouting an address from the first node to the second node, or by remapping a port of the first node to instead point to a port of the second node, in one of two different ways. Huang does not provide disclosure in which node failover is concerned at all. Rather, the type of network failure that Huang is concerned about, channel failover, while useful, does not pertain to node failover. Therefore, Huang cannot alone or in combination render the claimed invention unpatentable.

Richter relates to using different types of checksums to ensure that data split into packets is properly received at the other end. (See Abstract, for instance.) Richter can be used “in systems utilizing a distributed interconnect, such as for example, a switch fabric.” (Para. [0012]) But, splitting data into packets, and calculating checksums for such packets, does not pertain at all to node failover. For example, in the claimed invention, a node may send data in packets to another node. Utilizing Richter assists in ensuring that the data sent from the first node is properly received by the second node. However, Richter is not helpful in the situation where one of the nodes failed. By comparison, the claimed invention provides for failover of such a node in

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 23

accordance with two different types, path rerouting and port remapping, the latter of which the invention provides for two different types. Therefore, Richter cannot alone or in combination render the claimed invention unpatentable.

Wang is similar to Huang in its subject matter, by relating to the failover of a "link" or channel within a network. Thus, "[e]ach node has a cluster adapter connected to multiple port switches through communication *links*. . . . A fabric manager module will monitor the network and detect a link failure. Upon the detection of a link failure . . . a spanning tree partitioning module will partition the network into two trees at the point of the link failure. . . . The fabric manager module will then download the routing and distance table to only those switches effected [sic] by the new link selected to replace the failed link." (Abstract) By comparison, the claimed invention relates to failover of a *node*, not a *link* connecting to that node, as in Wang. In the claimed invention, a node fails over to another node, by rerouting an address from the first node to the second node, or by remapping a port of the first node to a port of the second node in one of two different ways. Wang does not provide disclosure in which node failover is concerned at all. Rather, like Huang, the type of network failure that Wang is concerned about, link failover, is useful but does not pertain to node failover. Therefore, Wang cannot alone or in combination render the claimed invention unpatentable.

Chen is probably one of the more relevant references cited by the Examiner, in that it deals with port spoofing. Port spoofing is only somewhat similar to the subject matter of the claimed invention, however. "[P]ort spoofing' . . . allows a computer to 'fail over' to its secondary fibrechannel connection if its primary fibrechannel connection should fail." (Col. 1, ll. 16-20) Note the difference between the type of failover that the port spoofing of Chen provides for compared to the type of failover that the port remapping (and path rerouting – which Chen does not mention at all) of the claimed invention provides. In port spoofing, if the primary fiber channel connection/port of a node (i.e., a computer) fails, the secondary fiber channel connection/port of the node can "spoof" the primary fiber channel connection/port. That is, what

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 24

is failing over is not a node to another node, as in the claimed invention, but rather a connection or port to another connection or port *of the same node*. Therefore, the type of failover that Chen is concerned with is useful and does deal with ports in some degree, but because Chen does not deal with node failover, but rather channel/port failover, and does not disclose path rerouting, Chen cannot alone or in combination render the claimed invention unpatentable.

Blumenau also relates to ports, but has nothing to do with any type of failover. Rather, Blumenau relates to using virtual ports to transfer data within a data storage system. Blumenau discloses interesting "virtual ports" that appear to hosts as "physical ports." (See Abstract.) Thus, Blumenau can allow a larger number of hosts to communicate with a data storage system, because a number of virtual ports can map to a single physical port of the data storage system. Blumenau does not discuss port remapping in either of the ways claimed in the claimed invention, however, even if it does pertain to ports generally and does not discuss node failover at all. Furthermore, Blumenau does not disclose path rerouting. For all of these reasons, Blumenau cannot alone or in combination render the claimed invention unpatentable.

Finally, Craft is similar to Blumenau in that it relates to ports, but has nothing to do with any type of failover. Rather, Craft relates to having intelligent network interface cards (NIC's) that allow a host computer to offload processing for multiple network connections. (See Abstract.) The ports for these multiple network connections are aggregated, so that to the host computer, only a single port is exposed. This is an interesting disclosure, but does not have anything to do with port remapping in either of the ways claimed in the claimed invention, and does not pertain to node failover at all. As with Blumenau, Craft also does not disclose path rerouting. For these reasons, Craft cannot alone or in combination render the claimed invention unpatentable.

Kashyap
Serial no. 09/917,314
Filed 7/27/2001
Attorney docket no. BEA920010015US1

Page 25

Conclusion

Applicants have made a diligent effort to place the pending claims in condition for allowance, and request that they so be allowed. However, should there remain unresolved issues that require adverse action, it is respectfully requested that the Examiner telephone Applicants' Attorney so that such issues may be resolved as expeditiously as possible. For these reasons, this application is now considered to be in condition for allowance and such action is earnestly solicited.

Respectfully Submitted,



March 4, 2005

Date

Michael A. Dryja, Reg. No. 39,662
Attorney/Agent for Applicant(s)

Law Offices of Michael Dryja
704 228th Ave NE #694
Sammamish, WA 98074
tel: 425-427-5094, fax 206-374-2819